

HIPPI-6400

Combining network innovations into a high throughput, low-latency, reliable networking standard.

HIPPI-6400 Features

- VLSI MAC chips available 2nd Q'97 with hardware deployment 1st Q'98.
- 6400 Mbit/s (800 Mbyte/s) full-duplex link-payload data rate.
- Scheduled Transfers enable DMA-engine to DMA-engine data movement.
- Flow controlled, reliable, hardware multiplexed link layer.
- Translation, not encapsulation, of legacy protocols to reduce destination processing overhead and promote interoperability.
- Striping support for HIPPI-6400 to HIPPI-800 translators.
- Pseudo RFC 1042 (Ethernet) MAC Header used for routing and transparent bridging to Ethernet.
- SNMP switch management with extreme code reuse from current Ethernet/IP designs.
- Direct support for upper layer protocols: IP, IPng (jumbograms), MPI, and IPI-3.
- No options (excluding copper or fiber physical interface) to ensure interoperability and keep design and installation simple.
- Local area transmission technology built for applications requiring 1 km maximum link distance.

D_ULA (msb)			
D_ULA		S_ULA (msb)	
S_ULA			
M_LEN			
802.2 LLC			
DSAP	SSAP	CTL	ORG
ORG	ORG	EtherType	
Opcode	Flags	S_Count	
R_Port		S_Port	
Key			
R_id		S_id	
Bufx			
Offset			
T_len			
B_num			
OS_Bufx			
OS_Offset			

Scheduled Transfer Header Format

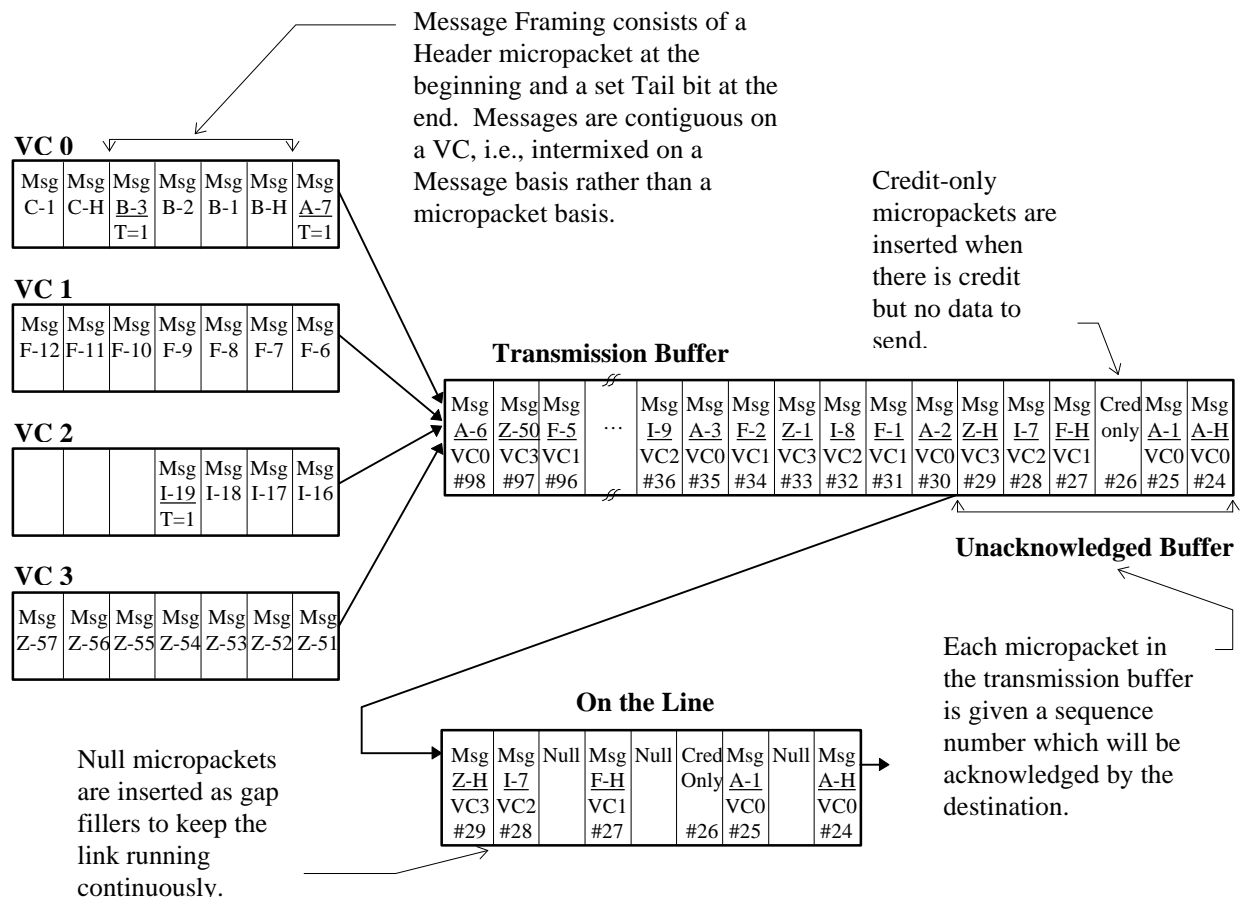
The HIPPI-6400 standard includes a Scheduled Transfer protocol that provides an interoperable method to predetermine the transfer size, buffer pools, buffer sizes, and ULP binding information, and to correctly tile into the destination's buffers. A flow-controlled transmission-unit hierarchy is provided that allows acknowledgment, retransmission, striping, and intelligent resource allocation. The Scheduled Transfer uses the virtual channel Message sizing conventions for separating control and data messages for greater concurrency and lower latency. The Scheduled Transfer allows implementors to create an operating system (and network driver) interface that can bypass the normal sources of software overhead to achieve transfer bandwidths limited only by the link capacity and DMA engine bandwidth.

The figure to the left shows the HIPPI-6400 MAC Header, followed by the 802.2 LLC/SNAP Header and the Scheduled Transfer parameters (64 bytes total). An 32 byte upper layer protocol payload may be sent after each Scheduled Transfer control operation to aid upper layer processing.

Logical Layer Example

HIPPI-6400 Messages are composed of micropackets which contain 32 data bytes and 8 control bytes. Each small box displayed in the figure below represents a micropacket. The first micropacket in a Message is the Header micropacket (the first 48 bytes in the top figure) and carries 48-bit source and destination ULA's for routing through switches. Each micropacket of a Message originates on one of four virtual channels (VC's) and travels throughout the fabric on the same VC. In the figure below each micropacket has a Message letter followed by the data micropacket number or H for the Header micropacket.

HIPPI-6400 ensures reliability by positively acknowledging micropackets and by retransmitting if an ACK is not received in time. HIPPI-6400 flow controls the link by issuing micropacket credits from VC receive buffers to the transmission state machine.



See "<http://www.cic-5.lanl.gov/~det/>" for more information including complete standards documents.